



Whistler Design Preview

April 20 & 21



Storage Snapshots for Backup

David Golds
Lead Program Manager,
Filesystems & Storage
Microsoft Corporation

One problem with backup

A program modifies a file while it is being backed-up...

File
contents
at time t_0 :

Blow up all the party balloons

.....

File
contents
at time t_1 :
Contents
of the
backup:

'Ship gift' for Win2000 team!

Blow up all the Win2000 team!

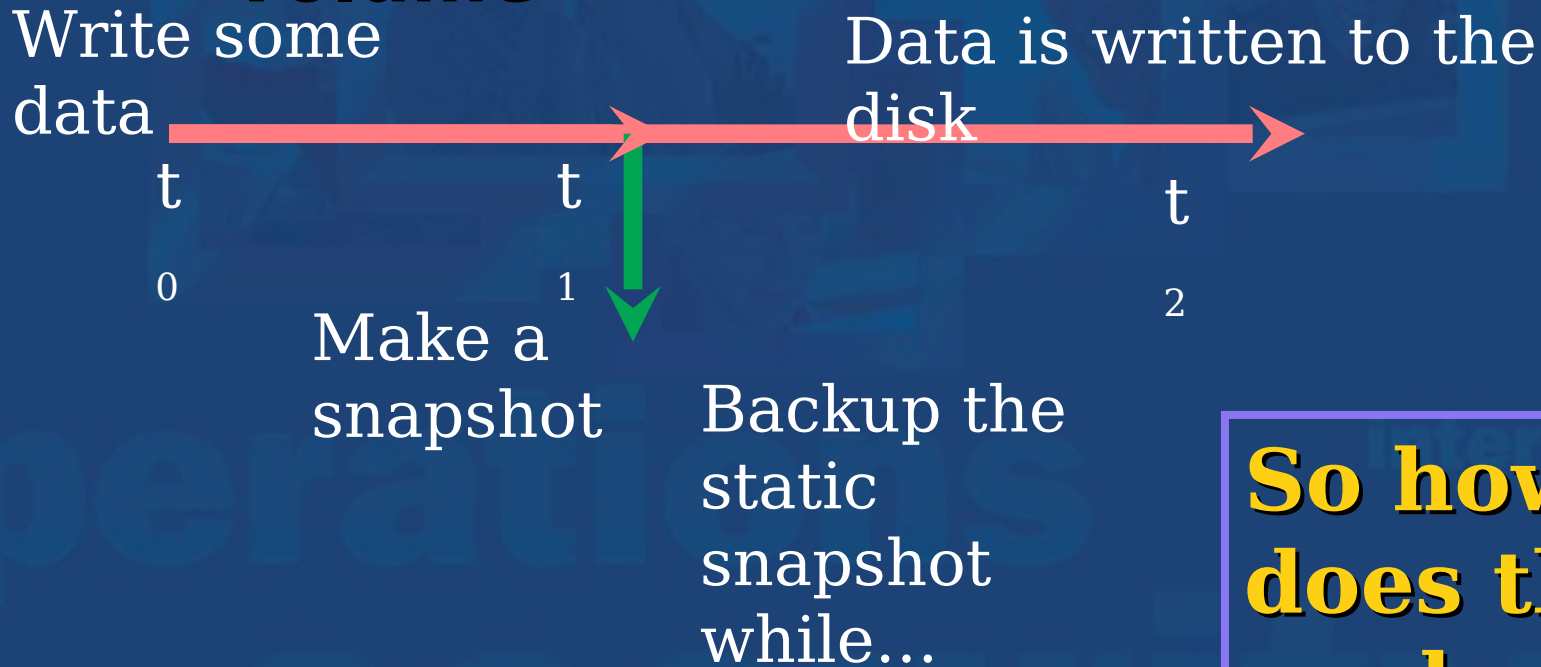
On my planet we call this 'data corruption'...

Four potential solutions

- Stop the applications while you do the backup
- Skip (exclusively) opened files during the backup
 - NT Backup does this. Sorry, did you want that file?
- Explicit APIs for each application
 - Elusive so far, Many different APIs in Win2000 backup for each service
- Take a 'snapshot' of the data, and back that up
 - Most 3rd party backup programs have an add-on for this
 - Use a variety of hardware or software schemes

What's a snapshot?

- A copy of some data set at a point in time
 - The data set here is a disk volume



**So how
does this
work...?**

Method #1: Split Mirrors

High End Solution

- Make a **complete copy** of the entire Volume
 - Step 1: Mirror the volume
 - Step 2: Break off one mirror
- Very fast once split
- Takes time to synch mirrors
- Expensive
 - Costs disk space.
 - Considered a 'High-End' storage feature

Method #2: Copy-on-write

Don't copy everything

- Only **copy what changes** since the snapshot
- Can be done as a (volume or file) filter or in hardware

Sectors on-disk at time of snapshot (t_1)	⟨Ĳ 2 Ž 4 5 “” •
Sectors modified between time t_1 and t_2	① ② ③ ④ ⑤ ⑥ ⑦ ⑧ ⑩
Copy replaced sectors to a side store:	⟨Ĳ Ž “” •
Recover old view (t_1) by laying snapshot 'side store' over current data	⟨Ĳ 2 Ž 4 5 “” •

This isn't rocket science. IBM was doing this 30 yrs ago.

What are we doing about it?

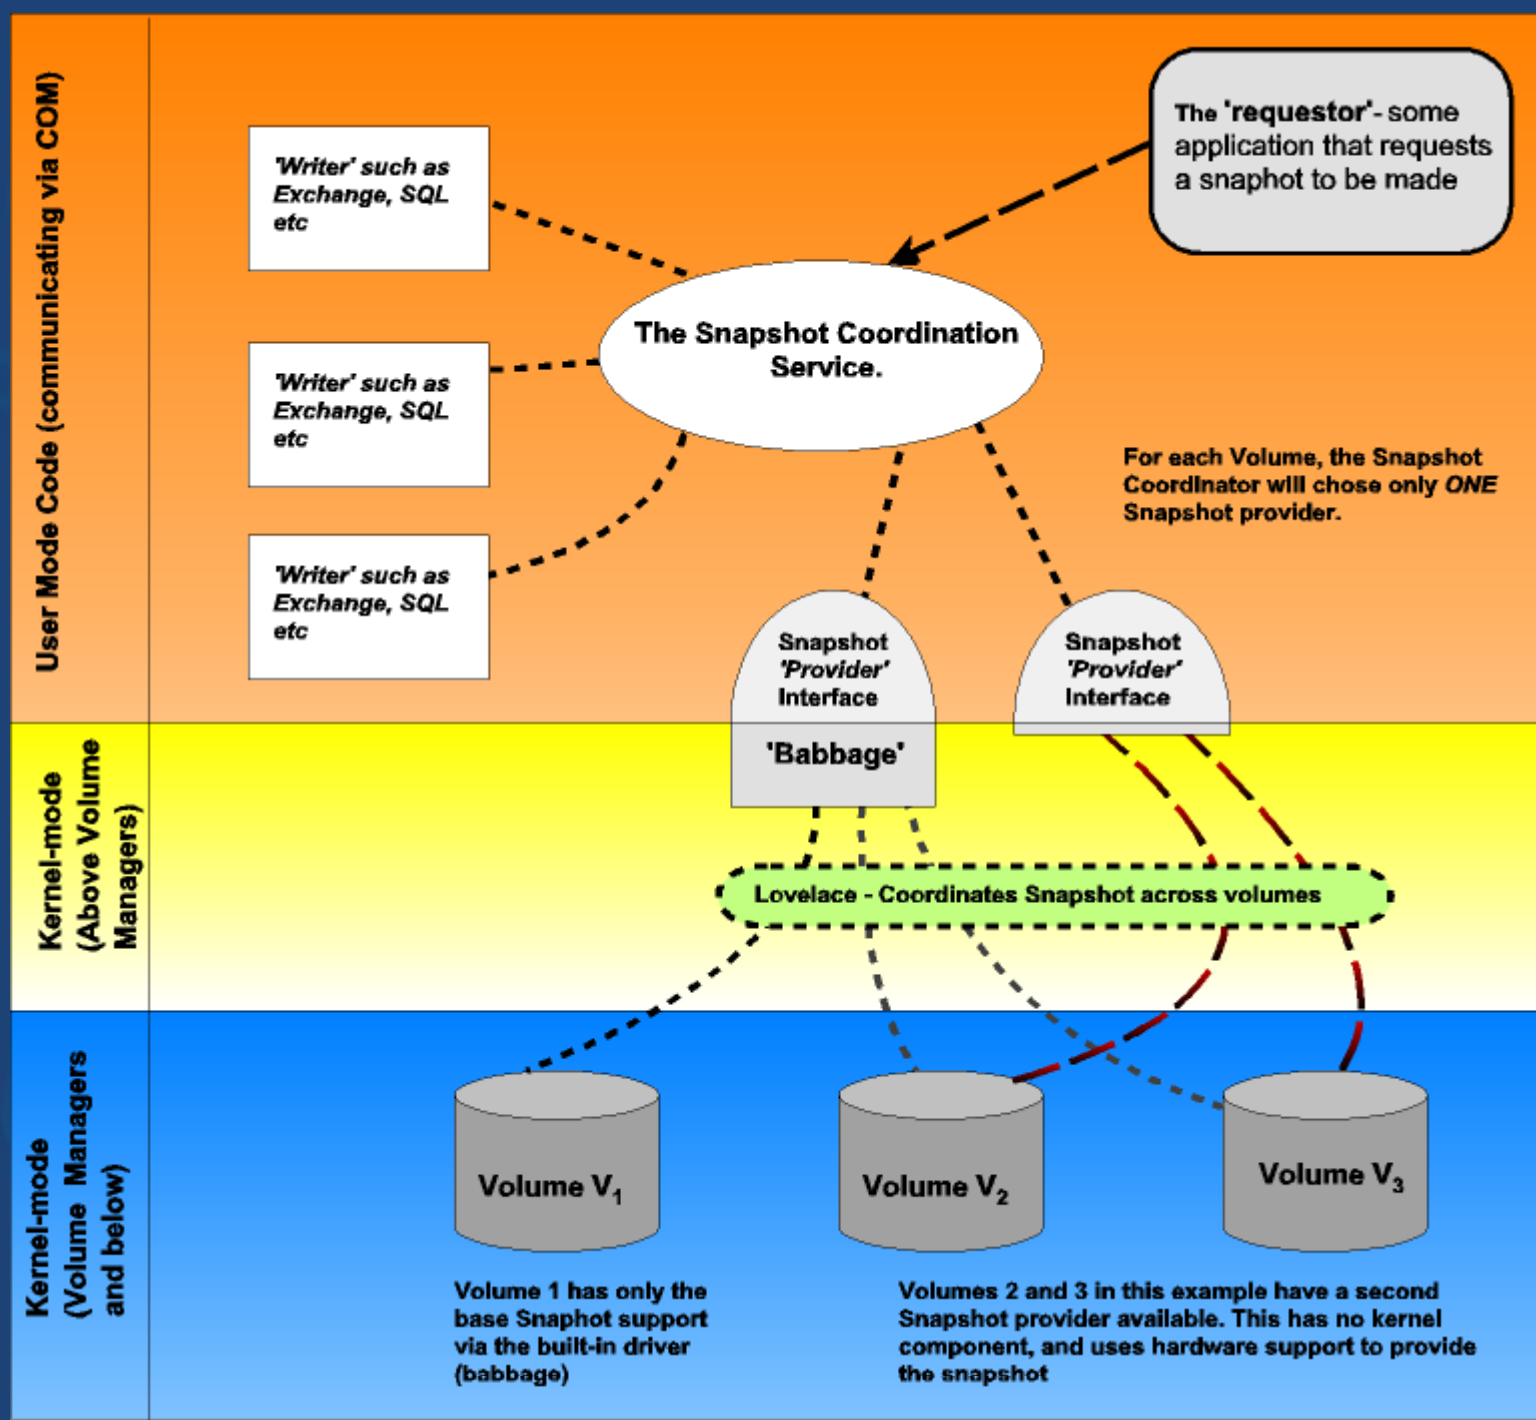
Make Windows a great platform for snapshots:

- **Provide a common API and coordinator to manage all of this.**
- **Help Microsoft applications and services, ISVs applications and stores and IHVs to adopt the interfaces**
- **Provide an in-box snapshot driver**
 - **Seed the market**
 - **Ensure concept generally available**

Snapshot design decisions

- Why did we choose snapshots as a general backup paradigm?
 - Even if the apps do nothing, you are still better off
 - It works at the low-end, via software drivers
 - It scales up to Enterprise, using Data Center hardware
 - It requires less effort for applications than many alternatives
- But still, some stores will

The Big picture



Snapshot Design Elements

- *Providers* do the actual snapshotting
- *Writers* are applications that are changing data
 - They get the easiest interface to write
- *The Volume Snapshot Service* just talks to **Writers and Providers**
- *Snapshot Sets* allows us to span volumes
- **Operations:**

Whistler deliverables (1 of 4)

- **Application and FS synchronization across volumes**
 - **Easy Freeze and Thaw APIs for 'Trained Applications'**
- We are evangelizing use of this API for MS and external applications
 - Any apps that hold files open
 - **If the App does not support this API?**
 - **All volumes snapped at a common point-in-time.**
- API also helps backup Apps find the data

Whistler deliverables (2 of 4)

- **Microsoft's Snapshot coordinator**
- **The coordinator has no persistent state**
 - **Snapshot information is held in the providers**
 - **The coordinator can aggregate queries**
 - **The coordinator finds and communicates with**
 - **Multiple Providers**
 - **Multiple Writers**
 - **Discovery and enumeration of snapshots**

Whistler deliverables (3 of 4)

- **Snapshot Provider interface**
 - **Enables IHVs and ISVs to plug in their snapshot engines**
 - **Numerous IHVs and ISVs will be plugging-in**
- **The in-box snapshot provider is designed to support the backup scenario only:**
 - **Read Only snapshots (NTFS can now mount RO)**
 - **Snapshots disappear upon reboot**
 - **Only snapshot one active on a volume at**

Whistler deliverables (4 of 4)

- **NtBackup will use snapshots**
 - **Snapshot multiple volumes**
 - **Reverts to non-snapshot if the snapshot fails**
- **But there's more..**
 - **The 'Writer Interface' can do more:**
 - **'Writer' can describe location of data**
 - **'Writer' can participate in Restore**

Where do you want to go today?®

Microsoft®